Redesigning Reservation: An Institutional Alternative to Quota-based reservation in India

Working Paper

Jayasuryan Iyer
jayasuryan_iyer@outlook.com
August 2026

Abstract

India's reservation system has never, in the record reviewed here, seen a benefit withdrawn once granted — a pattern usually treated as a political observation rather than argued as a mechanism. This paper makes two contributions. First, it states that pattern as a formal, testable diagnosis: group-based classification creates a rent that makes political mobilization the rational strategy for defending or expanding a category, while courts remain the only actor able to constrain the pattern, and only by setting an outer ceiling rather than redesigning the underlying allocation logic. Second, building on this diagnosis, it develops a fully specified institutional alternative — a continuous, multi-domain "deficit-point" system scoring verified individual circumstance rather than group membership, administered by a two-body governance structure designed with direct reference to the documented failure modes of India's own insulated institutions (the Reserve Bank, TRAI, SEBI, the Finance Commission, the Lokpal, and the National Judicial Appointments Commission) — a design that mitigates several of those failure modes but, as the paper states directly in Section 6.5 and Section 6.7, does not claim to close all of them, and leaves at least one (direct political pressure of the kind that produced the RBI's 2018 Section 7 episode) as an acknowledged, unresolved exposure. The comparative claim for individual over group targeting is stated in falsifiable form, with named conditions under which it should be expected to fail. The paper then turns its own diagnostic apparatus on its own proposal, concluding that no presently identifiable legal or political pathway would adopt this design under current conditions — and argues, following the ideal/non-ideal theory distinction in political philosophy, that this finding does not make the diagnosis or the design wrong, only currently unactionable.

Keywords: reservation, affirmative action, caste, India, institutional design, path dependency, delegation, non-majoritarian institutions

---

1. Introduction

India's reservation system was built to correct a specific historical wrong: centuries of caste-based exclusion from education, employment, and political power. That underlying wrong is real, and nothing in this paper argues otherwise. What this paper argues is narrower and, in a sense, more damning: the *mechanism* chosen to correct it — allocating benefits to entire castes and tribes rather than to individuals with verified need — has a structural flaw that has nothing to do with whether the original diagnosis was right. A system that scores groups creates strong incentives for expansion, because organized groups can mobilize around inclusion and protection while contraction imposes concentrated losses on established beneficiaries. Those incentives persist even when the underlying distribution of disadvantage changes.

The paper proceeds in eight sections. Section 2 situates the argument against the relevant literature and states, specifically, what is and is not original here. Section 3 sets out the case-selection logic and evidentiary approach. Section 4 lays out the evidence for the paper's first contribution: a formal, mechanism-based diagnosis of why India's reservation system expands but does not contract — a system with no documented instance, in the record reviewed here, of a caste, community, or benefit being removed once granted; whose only real constraint is the judiciary setting an outer ceiling it cannot itself redesign; and whose incentive structure rewards political mobilization over documented need. Section 5 asks a harder question implicit in that evidence: if legislatures reliably cave to organized pressure, is it even legitimate for government to resist doing so? The answer is yes, on the same general constitutional logic — insulating certain decisions from ordinary majority control when ordinary political incentives would otherwise distort them — that already governs several existing Indian institutions, and it explains why the alternative proposed here is deliberately insulated from ordinary political control rather than merely well-intentioned.

Section 6 is the paper's second contribution and its centerpiece: a full architecture for an individual-level "deficit-point" system — a continuous, verifiable, fraud-resistant scale that scores actual circumstances instead of group membership, governed by a two-body structure designed to be accountable without being politically capturable. It also states directly a claim the architecture alone does not make on its own: under what conditions individual scoring should be expected to outperform group classification, not merely operate differently from it. Section 7 then turns the paper's own diagnostic apparatus on the proposed design itself, asking whether it could actually be adopted, and concludes that no presently identifiable path would deliver it on a near-term timeline, for the same structural reasons diagnosed in Section 4. Section 8 addresses what that finding means for whether a proposal that is currently unadoptable is still worth making.

The claim running through all eight sections is that this is not a policy brief pitching a plan for near-term adoption. It is an argument about what a defensible system would look like, paired with a diagnosis of why no presently identifiable pathway exists for the current system to produce it. The two claims are kept deliberately separate: conflating a design's merits with its near-term political prospects is how sound diagnoses are discarded along with poor prospects for enacting them.

A scope limitation should be stated before the argument begins. Affirmative action can be justified on more than one ground: as compensation for a specific historical wrong, or as a present-oriented tool for allocating scarce opportunity to whoever is currently most disadvantaged. This paper does not claim that individual need is the only legitimate basis for affirmative action, or that historical, compensatory justice is an illegitimate goal. Its claim is narrower: if the objective is preferential allocation to those presently disadvantaged — which is how the existing system is in practice defended, not merely how it is theoretically justified — then group membership is an increasingly blunt instrument for identifying the recipient, and a system that measured disadvantage directly would better serve that same objective. A purely compensatory theory of what reservation is for is not addressed by this paper's case, and the paper does not claim to answer it.

---

2. Related Literature

Three bodies of scholarship bear directly on this paper's argument, and situating it against each clarifies what it adds.

Indian reservation politics. Galanter's account of how courts and legislatures jointly built India's classification system remains the reference point for the doctrinal history in Section 4. Deshpande and Ramachandran's 2024 analysis of 2019–21 National Family Health Survey data confirms that caste gaps in education, occupation, and wealth persist across generations, at rates that vary enough across the SC-ST/OBC/general split to complicate any single group-level account. Weisskopf's comparative work on U.S. and Indian affirmative action, and Jaffrelot's account of caste-based political mobilization, establish what this paper treats as background rather than a claim needing re-argument: that group classification functions as a mobilization mechanism, not only a targeting one.

Structural irreversibility: a formal argument for something only informally recognized. That reservation expands but never contracts is a common observation in Indian political commentary — editorials describe it as "an endless exercise in power-sharing," and the EWS quota's aftermath was widely labeled a "domino effect" within weeks of passage. What doesn't exist, in the academic literature searched here, is a formal mechanism-based argument for why: no application of Olson's concentrated-benefits/diffuse-costs logic, or Pierson's increasing-returns account of policy feedback, to Indian reservation specifically, and no systematic compilation of case evidence (Patidar, Jat, Maratha, Bihar, Telangana, EWS) tied to a documented-zero-rollbacks finding. Section 4 is this paper's first distinct contribution: taking a widely felt but never formally argued claim and stating it as a testable mechanism with compiled evidence, rather than editorial intuition. One tension is surfaced rather than avoided: Skocpol's related finding, that narrowly targeted benefits are politically weaker than broad-based ones precisely because they lack a mobilized constituency, predicts that this paper's own individually targeted design may be *more* vulnerable to erosion once built, not less — a real cost of the proposed architecture, addressed directly in Section 6.7 rather than assumed away.

Individual versus group targeting: the closest prior work, and how this paper differs from it. The core move this paper makes — score disadvantage on multiple factors rather than caste alone — is not new. Deshpande and Yadav proposed a four-principle alternative (evidence-based, multi-dimensional, attentive to intersection, attentive to degree of disadvantage) in 2006, with a worked disadvantage-point model and a passing suggestion for "something like a diversity and disadvantage commission" to administer it. Their model is scoped to higher-education admissions specifically — the disadvantage tables use school type and exam location as inputs, with no employment-side analogue — and their diagnosis targets *targeting inaccuracy within* the existing OBC quota (elite sub-caste capture, gender and urban/rural gaps), not the irreversibility of the classification system itself. A shorter 2025 treatment revives the same basic proposal without comparable institutional detail — that piece appeared in a venue (IJRASET) that appears on multiple standing predatory-journal registries (including the stand-alone-journal listing derived from Beall's list), so it is cited here only as evidence that the basic proposal continues to circulate, not relied on for any factual or empirical claim — and Teltumbde (2024), writing for civil-liberties audiences, has separately proposed a family-level points scheme spanning both education and employment, though without a worked governance mechanism. This paper's contribution relative to all three is architectural: a unified score across education and employment with continuous re-verification (Section 6.6), a fraud-resistance design specified domain by domain (Section 6.4), an evidentiary gate governing when new disadvantage categories are recognized (Section 6.2), and a comparative claim stated in falsifiable form (Section 6.7) — none of which appear in this literature, and none of which Deshpande and Yadav's proposal, predating EWS, the Rohini Commission's findings, and NJAC's strike-down by over a decade, could have built against the evidence base now available. Coate and Loury's finding that group quotas can depress a targeted group's own investment incentives is a distinct mechanism-level argument for individual targeting, developed in full in Section 6.7 rather than here, since it belongs alongside the comparative claim it supports rather than in this literature survey.

Delegation to non-majoritarian institutions. Majone's (2001) credible-commitment theory of insulated-body delegation is the direct theoretical ancestor of the Appointment Panel/Technical Body design in Section 6, and should be named as such rather than left as an implicit analogy to the RBI and Finance Commission. Recent work on regulatory bodies in drug-rationing (England and France) finds that insulation from politics doesn't reliably produce stability, and can make a delegated body a *site* of politicization rather than an escape from it — a finding that lands directly on this paper's own admission, already in Section 6.7, that formal insulation isn't immunity. No proposal found in Indian reservation-reform literature or adjacent think-tank output applies delegation theory to a reservation-administering body with this level of institutional detail

— that gap, alongside the irreversibility argument above, is where this paper's clearest contributions sit.

The closest institutional precedent, and why it is not a comparator. The nearest thing India has produced to a permanent, autonomous body for administering group disadvantage outside the reservation system itself is the Equal Opportunity Commission proposed by the Sachar Committee in 2006, alongside a companion diversity index and an autonomous assessment and monitoring authority. A draft EOC Bill existed by 2014 and was recirculated for inter-ministerial consultation as late as 2017; it has not, on the evidence available here, been enacted. The EOC is worth naming for two reasons that cut in opposite directions. First, it undercuts any claim that no one has proposed a permanent, statutorily autonomous body in this policy space: the idea itself is not new, and this paper should not imply otherwise. Second, the EOC differs from the architecture proposed in Section 6 on every dimension that matters to this paper's contribution: it is a discrimination-complaint and monitoring body modeled on the UK's Race Relations Act, not a mechanism for allocating individual education or employment benefits through continuous, verified scoring; its remit was socio-religious communities rather than caste-based reservation categories; and nothing in its available design documents specifies the kind of capture-resistant, two-body separation between appointment and substantive authority developed in Section 6.5, engineered against the particular documented failure modes of RBI, TRAI, and SEBI. The EOC's decade-long stall is also a second, independent data point for Section 7's diagnosis: a proposal with cabinet-level backing for an insulated body in adjacent policy territory failed to be enacted for reasons structurally similar to those argued there, alongside the Lokpal's own delay.

---

3. Methodology

This paper is a case study paired with a comparative institutional-design exercise, not a cross-national statistical test, and its case-selection logic follows from that.

3.1 Why India

India is treated here as a most-likely case for the mechanism this paper diagnoses: a constitutionally entrenched, decades-old, group-based classification system operating under a competitive multi-party democracy with strong incentives for organized political mobilization around category boundaries. If the concentrated-benefits/diffuse-costs dynamic described in Section 4 does not produce the documented pattern of expansion-without-contraction in a system with India's scale, duration, and political salience, it is unlikely to be a general feature of group-based classification at all. The selection is deliberate rather than a matter of convenience or of the author's familiarity with the jurisdiction, and it carries a corresponding limitation, stated here rather than left for a reviewer to raise: a most-likely case is chosen precisely because it is expected to show the mechanism strongly, which means the findings do not by themselves establish how the same dynamic would play out in newer, smaller, or less politically entrenched quota systems, where the same incentive structure may not yet have had time to produce comparable path dependency. The paper's claims about mechanism are argued to generalize in principle — the underlying logic is Olson's and Pierson's, not specific to India — but the *evidence* for how the mechanism actually plays out is drawn from one case, and the paper does not claim more than that. This generalization claim should not be read as extending to the governance architecture proposed in Section 6: that design is deliberately built against the documented failure record of specific Indian institutions (RBI, TRAI, SEBI), and its capture-resistance features are correspondingly tuned to India's own regulatory history rather than to a general theory of insulated-body design. A jurisdiction with a different institutional failure record would need the diagnosis in Section 4 re-argued on its own evidence and the architecture in Section 6 substantially re-derived, not merely re-labeled; only the former is claimed to generalize as stated, and the latter is not.

3.2 Why these comparators, and not others

Section 6's governance architecture draws on a specific set of Indian institutions — the Reserve Bank of India, the Telecom Regulatory Authority of India, the Securities and Exchange Board of India, the Finance Commission, the Comptroller and Auditor General, the GST Council, the Lokpal, and the National Judicial Appointments Commission — plus one non-Indian comparator, the United Kingdom's contextual-admissions system and its Cardiff-model scoring instrument. These were selected on a single shared criterion: each is a documented, real-world attempt to insulate a specific administrative function from ordinary majoritarian control, with enough of a public record — statutory text, litigation, or reported institutional history — to identify what specifically succeeded or failed and why. The Indian institutions were prioritized over foreign ones for the governance-design comparisons specifically because a design meant for adoption within India's own constitutional and administrative structure should be built, where possible, against failure modes documented in India's own institutions rather than imported ones; a domestic precedent is harder for a domestic reviewer to dismiss as inapplicable. The UK case is the one deliberate exception, included not as a governance-insulation comparator but for a narrower purpose: it is the closest operating example of a continuous, individual-level disadvantage-scoring instrument at real administrative scale, which the Indian comparators do not provide.

This selection leaves out systems that are, in some respects, closer analogues to Indian reservation than any comparator actually used here. Malaysia's bumiputera policy and Brazil's racial quota system are both group-based affirmative-action regimes facing criticisms of entrenchment and mobilization-driven expansion that parallel this paper's diagnosis of India more closely than the UK's individually-targeted contextual-admissions system does. They are not treated as comparators in this paper, and that omission is a scoping choice rather than an oversight: this paper's second contribution is a governance-architecture proposal, not a comparative theory of group-based affirmative-action politics across countries, and the institutions selected for Section 6 were chosen for what they show about insulating an administrative body from majoritarian capture, not for how closely their underlying policy resembles reservation. A comparative treatment of Malaysia and Brazil alongside India would be a different, and defensible, paper; it is not this one, and a future-research note to that effect appears in Section 8.

3.3 Evidentiary approach

The evidence base is legal and documentary case analysis — constitutional text, statutes, Supreme Court and High Court judgments, government commission reports, and institutional records — supplemented by secondary quantitative literature (survey-based studies of caste and wealth inequality, targeting-accuracy research from the development-economics literature, and existing empirical work on the specific institutions discussed in Section 6). This paper does not collect original data of any kind: no survey, no interview material, and no original statistical analysis. Where a quantitative claim appears — the Rohini Commission's sub-caste representation findings, the Sutton Trust's correlation figures for POLAR, the wealth-variance findings from National Sample Survey data — it is drawn from a published source and cited as such, not derived independently. This is a deliberate scope limitation as much as a description of method: a paper making an institutional-design argument of this scale does not need, and could not responsibly attempt within its scope, an independent empirical test of every quantitative claim it relies on secondhand.

3.4 Genre

This paper mixes three modes of argument that are conventionally kept separate, and does so deliberately rather than by drift: a normative argument (Section 5, on whether insulating this kind of decision from majoritarian control is legitimate), an institutional-design proposal (Section 6, specifying a governance architecture in enough operational detail to be evaluated, not merely gestured at), and a political-economy diagnosis (Sections 4 and 7, arguing a mechanism-based account of why the current system behaves as it does and why the proposed alternative is unlikely to be adopted under present conditions). None of the three modes is subordinate to the others, and the paper does not resolve into a single genre by its conclusion — Section 8 argues explicitly that holding the normative and the diagnostic claims apart, rather than collapsing one into the other, is itself part of the paper's contribution. This paper does not deliver a purely empirical test of a hypothesis, nor a purely normative argument for a policy position, in isolation; that expectation is set explicitly here, rather than left implicit, before Section 4 begins.

---

4. The Case Against the Current System

India's reservation system is not simply unfair in its details; on the evidence reviewed here, it lacks a credible endogenous mechanism for systematic correction. Three failures compound: benefits granted have not been shown in the evidence reviewed here to be voluntarily withdrawn once granted, the political incentives created by group-based classification reward fragmentation rather than reduce it, and the design was never built to target actual individual need in the first place. None of this is a claim that redressing caste-based disadvantage is wrong. It is a claim that the mechanism chosen to redress it has a shape that creates strong incentives for expansion and makes retrenchment unusually difficult, regardless of whether expansion still serves the disadvantage it was built to address.

4.1 What the system actually is

Reservation in India is caste- and tribe-based, not religion-based; religious quotas have been tried and struck down by courts, including a 2010 ruling against an Andhra Pradesh scheme reserving seats for Muslims. The system does contain a real double-dip mechanism: Tamil Nadu's Backward Class framework includes a 3.5 percent sub-quota for Backward Class Muslims — a caste-linked benefit tied to religious community — which can stack with separate scholarship schemes run by the central government's Ministry of Minority Affairs for religious minorities generally. Two distinct programs, drawing on the same underlying marker, without checking against each other.

The creamy layer — the income-based exclusion meant to keep the wealthiest members of a reserved category from claiming benefits meant for the disadvantaged within it — was not part of the original design. The government implemented a flat 27 percent quota for Other Backward Classes in 1990 with no income filtering at all; the Supreme Court retrofitted the creamy layer test two years later, in *Indra Sawhney* (1992). It still does not apply consistently to Scheduled Castes and Scheduled Tribes — a question the Court left unsettled again in *State of Punjab v. Davinder Singh* (2024) — and even its application to OBCs remains contested, as shown when Haryana's 2016 attempt to define it purely by income was struck down in 2021. A patch bolted onto a system after the fact, applied unevenly across the categories it was meant to govern, is not a design; it is a running repair.

The framers' choice of caste over economic criteria in 1948–49 was deliberate, not an oversight born of inexperience — but the mechanism built alongside it to keep that choice under review turns out to have been weaker than it looks. Article 16, covering employment reservation, was in the Constitution's original draft. Article 15, covering education reservation, was not: it was added reactively through the First Amendment in 1951, after courts struck down an earlier scheme in Madras, and it received no sunset clause at all. Political reservation under Article 334 did receive a ten-year sunset — the one accountability mechanism the framers built into the system — but it has been renewed seven times over more than eighty years, and the historical record complicates how much of a real check it was ever meant to be. Scholarship drawing on the Constituent Assembly debates (Bhaskar 2022) indicates the ten-year limit was a compromise imposed by other members over the preference of B.R. Ambedkar, the policy's principal architect, who favored extending political reservation through repeated constitutional amendment rather than allowing any sunset to lapse. That the claim needs correcting matters beyond historiography: Justice Pardiwala repeated the standard "Ambedkar wanted only ten years" account from the bench in the 2022 EWS case, *Janhit Abhiyan* (2022) — meaning the myth this scholarship rebuts had reached the Supreme Court itself, not just popular retelling. The narrower historical point is therefore that the often-repeated ten-year-sunset story is incomplete: the constitutional record shows a limited accountability mechanism whose operation was more complicated than the standard account suggests. This paper does not need to resolve every question about Ambedkar's personal preference; it needs only the institutional conclusion that the original architecture did not provide a robust self-correction mechanism.

4.2 The pattern: courts constrain, legislatures rarely do

This paper has found no documented instance, in the evidence reviewed here, of a caste or community being removed from the Scheduled Caste, Scheduled Tribe, or Other Backward Class lists once added, or of a legislature voluntarily narrowing a reservation scheme it had already granted. Every documented attempt to constrain reservation's growth in that same evidence has come from courts, and the interventions reviewed here have repeatedly been resisted, relitigated, or worked around rather than accepted as durable stopping points. The Maratha quota — a politically dominant caste, in the Supreme Court's own words "in the mainstream of national life" — was struck down in 2021 (Patil v. Chief Minister 2021) and re-attempted three times across a decade: 2018, 2019, and 2024. Bihar's attempt to push reservation to 65 percent was struck down by the Patna High Court in 2024. Telangana's attempt to raise a local-body quota from roughly 25 percent to 42 percent for Backward Classes, pushing the total to 67 percent, was stayed by the Telangana High Court in October 2025, a stay the Supreme Court declined to disturb days later, reaffirming the 50 percent ceiling first set in *Indra Sawhney v. Union of India* (1992).

The absence of a located counterexample is a claim about the evidence reviewed, not a claim to have exhaustively searched every state and every classification revision in India's history, and the distinction is maintained deliberately. The theoretical explanation for why the pattern should exist regardless is a separate claim, developed below, and the two are kept apart deliberately: the empirical pattern and the reason to expect it are different kinds of evidence, and treating them as one blurs how much weight either can bear alone.

The Economically Weaker Sections (EWS) quota, passed in 2019, is the clearest illustration of how thin the legislative resistance actually is. It cleared the Lok Sabha 323 to 3 on 8 January, and the Rajya Sabha 165 to 7 the following day — both houses within a single week of the bill's introduction. An opposition member of parliament, Kapil Sibal, put the political calculation on the record at the time: opposing the measure would cost more than supporting it would help. The move followed the ruling party's loss of three state elections and has been read as a response aimed at an upper-caste voting bloc. The pattern that followed was immediate rather than gradual: within weeks, Gujarat had restricted its own EWS eligibility criteria, Andhra Pradesh had carved out a subcaste share of its own, and a state party leader in Bihar was publicly demanding 90 percent reservation. A single new category, once granted, can alter the political expectations of other groups about what is achievable.

4.3 The mechanism, and why the examples that follow illustrate it rather than prove it alone

The incentive built into a group-based system is not incidental to its failures; it is the mechanism producing them, and stating it precisely before turning to examples matters because the examples can otherwise be read as proving more than they do. The claim is not that reservation policy manufactures caste identity or caste-based political organization from nothing. Caste identity in India long predates, and would persist independently of, any specific reservation scheme, and historically salient groups would likely organize politically around *something* regardless of what benefits are on offer. The narrower and better-supported claim is this: attaching a scarce, durable material benefit specifically to *group classification* creates a distinct incentive to contest classification as such — to seek inclusion in, or to defend, a particular category boundary — that would not exist, or would exist in a much weaker form, if the same benefit were instead allocated on verified individual circumstances. Reservation does not create caste politics; it creates a specific reason for caste politics to organize around the classification boundary rather than around something else.

The evidence for that narrower claim is not abstract. Marathas — a dominant caste by most economic measures — have been repeatedly denied reservation status and have repeatedly remobilized to seek it. Patidars in Gujarat organized a major agitation in 2015 demanding OBC status. Jats in Haryana did the same in 2016, in an agitation that turned violent and killed dozens. In Rajasthan, Gujjars and Meenas have been locked in direct competition over Scheduled Tribe versus Other Backward Class status — a genuinely zero-sum conflict between two groups, neither of which can advance without displacing the other's claim. None of these are disputes over whether disadvantage exists; they are disputes over classification, which is exactly the kind of conflict a system that makes classification the gateway to benefit will manufacture, regardless of how much organizational capacity the groups involved would have had for other purposes.

This dynamic fits a broader pattern described by Mancur Olson (1965): concentrated benefits and diffuse costs favor organized collective action by the group that stands to gain. It cuts in the reverse direction, too — the same logic that explains why existing groups mobilize to expand or defend their status also suggests that an individual-based system would face less organized resistance to being narrowed later, since there would be no natural bloc with a shared interest in defending a specific classification. That advantage is not unqualified. Theda Skocpol's (1991) argument that narrowly targeted, means-tested benefits are politically weaker than broad-based ones — often summarized as "programs for the poor tend to be poor programs," though this paper attributes the underlying claim rather than a verified verbatim quotation — holds that diffuse, means-tested benefits are more politically vulnerable to underfunding and neglect than benefits organized around an identifiable, mobilized constituency, and is a real counter-consideration, not a dismissible one. Paul Pierson's (1994) research on welfare-state retrenchment complicates it further in the other direction: voters tend to punish visible losses more than they reward savings, which means diffuse programs are not defenseless so much as defended reactively, through blame-avoidance politics, rather than expanded proactively the way concentrated-benefit programs are. An individual-based system would not escape political vulnerability; it would trade one kind of vulnerability — capture and irreversible expansion — for a different kind — neglect and underfunding.

4.4 Naming the mechanism

The diagnosis above is stated here as a single named sequence, both because it is this paper's most original claim and because naming it makes it easier to check against evidence rather than accept on the strength of the examples alone: (1) government defines a group; (2) group membership determines access to a scarce benefit; (3) the benefit creates a rent attached to classification; (4) the rent makes classification politically valuable; (5) political mobilization seeks to gain or defend that classification; (6) successful mobilization expands the classified population or the benefit's scope; (7) expansion makes contraction politically costly; (8) the classification becomes institutionally persistent even as the underlying disadvantage it was meant to track changes, narrows, or shifts to different people. Call this the classification-rent mechanism. Section 6's central design choice — replacing group membership with individually verified circumstance as the trigger for benefit — is best understood as an attempt to break the sequence at step two, before rent, mobilization, or entrenchment can accumulate around a durable classification at all.

4.5 What this section establishes, and what it doesn't

None of this amounts to an argument that caste-based disadvantage is not real, or that redress is unwarranted. It is an argument about mechanism: a system that allocates benefits to groups rather than to verified individual circumstances creates strong incentives for expansion, because organized groups can mobilize around inclusion and protection, while contraction imposes concentrated losses on established beneficiaries, and courts — the only actor shown to constrain this pattern at all — can set an outer ceiling but cannot redesign the underlying allocation logic. What would need to replace it to do better is the subject of Section 6. First, though, a prior question: is resisting a mobilized majority's demands even a legitimate thing for government to do?

---

5. Why This Is Indefensible, Not Just Unfair

Section 4 is a positive account: a description of what tends to happen. Legislatures expand group-based benefits and rarely retract them; courts are the only actor that reliably constrains that expansion, and even they only set an outer ceiling rather than redesigning the underlying logic. Read on its own, that account risks being mistaken for an argument that government should simply do whatever a large, well-organized group demands, since that is what the evidence shows tends to occur. It is not that argument, and the distinction is stated explicitly here before turning to what should replace the current system.

5.1 The Madisonian starting point

The danger of an organized faction capturing the state to the detriment of others is not a new observation. It is the founding problem of constitutional design itself. Madison's argument in *Federalist* No. 10 (1787) is that pure majoritarian democracy is dangerous precisely because an organized faction — even a majority one — can use ordinary democratic process to harm the diffuse or the minority. The remedy he proposes is not to abandon democracy but to structure it: separated powers, checks between institutions, entrenched rights, so that organizing a majority is not, by itself, sufficient to run over what should not be run over.

India's own constitutional architecture reflects the identical logic, independently of anything borrowed from Madison. Part III of the Constitution — fundamental rights — exists because the framers did not trust a bare parliamentary majority to police itself on rights questions. Judicial review, and later the basic-structure doctrine established in *Kesavananda Bharati* (1973), exist for the same reason at a higher level: even a constitutional-amendment-strength majority cannot alter certain things, precisely because a large enough organized coalition getting its way is not treated as self-legitimating.

5.2 Why reservation is a different kind of question, and what generalizes from the fundamental-rights case

The bright line separating rights questions from ordinary politics is useful precisely because it is a bright line: there is no serious contestation, across constitutional traditions, courts, or international law, about certain categories of harm being placed outside majoritarian bargaining. Reservation design — how much redistribution, targeted by what criteria, phased on what timeline — is not that kind of question. It is a genuinely contested distributive-policy question, one on which reasonable constitutions, courts, and democratic majorities differ, and differ across time within the same country.

That distinction matters for what follows, because the argument for insulating the design proposed in Section 6 should not lean on an analogy between reservation and fundamental rights — the two are not equivalent in kind, and treating them as though they were invites exactly the objection that this paper is quietly upgrading a distributive-policy preference into a rights claim it hasn't earned. The narrower and more defensible version of the argument is this: the *general design principle* underlying fundamental-rights insulation — that some decisions warrant structural protection from ordinary majority control specifically because ordinary political incentives would otherwise systematically distort them — can extend to distributive administration without treating the underlying distributive choice itself as a fundamental right. Courts, the Reserve Bank, the Election Commission, and the Finance Commission are all institutions India already accepts as legitimately insulated from day-to-day majority control on specific technical or rights-adjacent questions, not because those questions are all fundamental rights, but because full majoritarian control of those functions was judged worse than the accountability cost of insulating them. That is the precedent this paper actually needs, and it is a lower bar to clear than the fundamental-rights analogy implies.

This distinction is not asserted here in the abstract; Section 4's own evidence already shows a version of it operating in practice. Courts do intervene on reservation — the fifty percent cap from *Indra Sawhney* (1992), the Maratha quota struck down (Patil v. Chief Minister 2021), Bihar's sixty-five percent struck down (Patna High Court 2024), Telangana's local-body quota stayed (Telangana High Court 2025). That is evidence of a counter-majoritarian check functioning in these cases. What it is not doing, and this is the actual gap the rest of this paper responds to, is improving the system's design. Courts can set an outer ceiling and police procedural overreach; the evidence reviewed here does not establish that they can or will redesign the underlying allocation logic toward something like individual-need targeting. A court can say "no further" and "not like this." It cannot say "here is a better system" and make it so.

5.3 What this means for the design that follows

The governance structure proposed in the next section — an appointed body insulated from ordinary parliamentary majorities, empowered to set binding rather than merely advisory policy — is an attempt to build a second, purpose-built check for exactly this class of question, on the design principle just described, where the check that already exists is real but structurally limited to ceiling-setting rather than active redesign.

The binding authority itself needs a narrower scope than it might first appear to claim, and that scope is stated precisely because "binding" is the word most likely to draw objection. What is proposed as binding is *methodology* — how disadvantage is measured, how factors are weighted, how verification is conducted — not the distributive outcome itself in the sense of a fixed allocation Parliament cannot revisit. Parliament retains the power to set the statutory objective, the constitutional boundaries, and, if it chooses, an aggregate ceiling on preferential allocation; the Technical Body's binding authority operates only within that mandate, on the technical questions of how the objective is measured and pursued, not on whether the objective itself should exist. That is a meaningfully narrower claim than a technical body determining, without review, who gets preferential access to education and employment — and it is the version of "binding" this paper can actually defend, rather than a broader one dressed in the same word.

It also explains a design choice that might otherwise look arbitrary: giving the proposed technical body's weighting recommendations binding force by default, overridable only by legislative supermajority with public justification. That is the calibration made concrete — insulated enough that a simple majority cannot casually overrule inconvenient findings the way it currently can override the absence of any real constraint on reservation policy, but not so insulated that it escapes democratic accountability altogether.

5.4 The epistemic objection to no reserved seat, not just the political one

The same logic underlies rejecting a reserved caste seat on the Technical Body in favor of qualification criteria and process transparency instead — but that design choice faces two different objections, and only one of them is answered by what has been said so far. The first, political, objection is that a reserved seat would reproduce the exact capture mechanism this proposal exists to escape; the response to that objection is the one already given, and it holds.

The second objection is sharper and different in kind: not that a reserved seat would be *captured*, but that someone without lived experience of caste-based discrimination may simply be missing information that a statistical model, and a qualification-filtered technical seat, cannot fully substitute for — pattern recognition of disadvantage that doesn't show up cleanly in administrative data, or an understanding of how a specific form of exclusion actually operates day to day. This is an epistemic claim, not a political one, and it deserves its own answer rather than being folded into the capture defense, where it does not actually get addressed. The design's answer is that lived and field familiarity with caste-based disadvantage is written into the social-stratification seat's qualification criteria as a substantive requirement, not merely permitted — the seat is not open to a generic statistician, but specifically to someone with demonstrated expertise in caste-based social stratification, which in practice will very often mean someone who has that lived or closely observed familiarity. What is rejected is not the *knowledge*, only its attachment to a fixed, singular, group-linked seat that cannot be revoked or contested and that would need, once created, an answer to the question of which single group's lived experience is privileged enough to reserve a seat for it — a question with no principled stopping point once opened. Qualification without reservation keeps the knowledge requirement while avoiding the durable-classification problem Section 4 diagnoses. Whether that resolution fully satisfies the epistemic objection in its strongest form is a question this paper does not claim to have closed; it claims only to have engaged the objection directly rather than treated it as answered by the capture argument alone.

A sharper version of this objection targets appointment, not qualification: an unreserved Appointment Panel can still produce a run of social-stratification appointees drawn from elite academic backgrounds, whatever the job description says. The fix is procedural, not compositional: before the Panel votes on the Search Committee's shortlist for this seat, the shortlist and its rationale are published for a fixed comment period — advisory only, no veto, no binding claim on the seat, so the reserved-seat problem doesn't re-enter through the back door. It puts the appointment's credibility on record at the point the objection actually bites, instead of resting on the criteria's wording alone.

5.5 What this does not establish

Establishing that the goal is legitimate to pursue does not make it easier to achieve. The political-economy evidence already presented — the absence of any precedent, in this paper's own review, for a legislature voluntarily narrowing an organized group's reservation benefit, and the incentive structure that rewards mobilization over documentation — is a claim about what tends to happen, and normative legitimacy does not change empirical difficulty. A constitutional design can be entirely correct in principle and still have no realistic path to enactment against a large, organized, electorally load-bearing constituency. Those are separate findings, and this paper keeps them separate rather than letting the argument in this section imply that the problem of adoption has somehow been solved along the way.

5.6 The technocracy objection

A natural objection is that insulating a technical body from majoritarian control simply reproduces, in a different form, the unaccountable elite capture that reservation's own defenders would say the system exists to correct — that this is technocracy dressed in constitutional language. The distinction is one of calibration, not category, and is the same question every insulated public institution has to answer; it is not evidence that insulation itself is illegitimate. A binding-but-supermajority-overridable structure, transparency and consultation requirements, and a removal mechanism that separates political accountability from the final decision to remove someone are the specific devices meant to keep a new technical body on the Finance Commission side of that line rather than drifting toward unaccountable technocracy. Whether they succeed is a design question taken up directly in Section 6.

---

6. What Should Replace It: The Deficit-Point System

If Section 5 established that India's reservation system fails because it scores groups rather than people, the obligation is to show what scoring people would actually look like — not as a slogan, but as a working architecture: a scale, a set of factors, a way of verifying them, a body to administer them, and rules for what happens when someone's circumstances change. Two existing systems are useful negative references throughout, because each shows a distinct way an individual-need system can fail even when its intentions are sound.

The first is India's own EWS quota. It uses a single flat income cutoff — ₹8 lakh, uniform regardless of whether the applicant lives in rural Bihar or urban Mumbai — a gap the Supreme Court itself pressed the government on (Chandrachud 2021) without receiving a socioeconomic justification. Eligibility relies substantially on applicant documentation and certification, and reporting cited here describes inconsistent field verification and limited systematic cross-checking against tax, land, and banking records (Counterview 2026). That reporting also identifies EWS-quota candidates with access to elite schooling and costly examination coaching (Indian Masterminds 2026), illustrating the risk that a broad eligibility rule can admit applicants whose circumstances differ substantially from the disadvantage the policy seeks to target. And because the cutoff is a single yes/no line rather than a continuous measure, a family earning ₹7.9 lakh and a family earning ₹1 lakh receive an identical benefit, while a family at ₹8.1 lakh receives none. That discontinuity has no relationship to the underlying need it claims to measure.

The second is the United Kingdom's contextual-admissions system, specifically its use of POLAR — an area-level deprivation index used as a proxy for individual family circumstance. POLAR captures income deprivation only weakly (a correlation of roughly 0.22, per Jerrim (2021); its overall correlation with family income is somewhat higher but still modest), which makes it an imperfect proxy for individual circumstance and creates both false positives and false negatives. This case is treated here as a lesson about proxy construction generally rather than as a direct analogue to Indian reservation — the two systems serve different populations under different constitutional and administrative constraints, and little is gained by treating British contextual admissions and Indian reservation as directly comparable policy contexts. What transfers, independent of that comparability question, is narrower: area-level data cannot substitute for individually verified circumstance, and any design that leans on it as a primary input inherits its weakness.

A workable system has to address both failure modes explicitly, not merely be careful about them at the point of calibration, since calibration is exactly where discretion erodes.

6.1 The scale: continuous, additive, bounded

The proposal adapts the model already operating at real administrative scale in UK contextual admissions — Cardiff's 0–330 additive scale, which sums weighted contributions from verified factors rather than applying a single cutoff. A continuous scale removes the arbitrary, off-target discontinuity by construction: moving from ₹7.9 lakh to ₹8.1 lakh in family income moves a score by a small, proportionate amount, not to zero. It also gives whatever body administers the system a natural lever for revision — weights can be nudged incrementally in response to outcome data, rather than requiring a wholesale eligibility redesign each time a threshold looks miscalibrated. It does not remove the fact that scarce slots produce a rank-ordered cutoff at the margin; no allocation mechanism under scarcity can, and the comparative claim this design can actually support is developed directly in Section 6.7.

The scale is additive *and* offsetting: some factors reduce a score as well as raise it. A documented reduction in disadvantage — inherited property, a family's income recovering — offsets accumulated points the same way accumulated disadvantage adds them. The score has to remain a live reflection of current circumstances, not a ratchet that moves in only one direction.

6.2 The factor domains

Four domains are scored directly. A fifth category — new forms of disadvantage not yet recognized by the system — is handled through a governance interface described below, rather than folded in by default.

Domain A: caste, tribal disadvantage, and gender. This is not a binary group flag but a verified individual marker, drawing on the certification infrastructure India already has — caste certificates, ST community-acceptance verification — but feeding into the continuous scale as one weighted input among several, rather than operating as a separate parallel quota. Gender belongs in this domain on the same logic, and has real doctrinal footing: *NALSA v. Union of India* (2014) recognized third-gender status and directed OBC-comparable treatment, and Article 243D already reserves panchayat seats for women. This is also the domain most exposed to fraud — fake and forged caste certificates are a well-documented, recurring problem — which is why verification here layers new audit machinery onto the existing scrutiny-committee system rather than replacing it. How well that layered verification actually works bears directly on how much this design can claim to improve on group classification at all, a question taken up in full in Section 6.7.

Domain B: financial circumstances. This domain is the direct answer to EWS's specific, documented failures. Income is measured continuously, not as a cutoff. It is regionally adjusted for cost of living: the rural/urban comparability problem raised in EWS litigation. The cited exchange did not establish a socioeconomic justification for the uniform threshold. And it is verified from multiple independent sources — PAN-linked income tax records, GST filings for business income, land records for agricultural holdings — rather than a single self-declared certificate. Assets are scored as a continuous offsetting factor rather than EWS's blunt land-acreage and flat-size cutoffs, which the government's own expert committee eventually dropped as too crude to reflect actual economic condition.

Domain C: family and household circumstances. This domain follows the strongest precedent identified in the UK system: free-school-meal eligibility, a direct individually verified indicator, rather than a broad income band or an area proxy. Candidate factors include single-parent household status, disability within the immediate family, and orphan or ward status. One boundary has to be enforced strictly: Domain C scores only the *non-monetary* consequences of a household event, never its income effect, which Domain B already owns. A job loss is scored in Domain C only for its disruption to childcare or schooling continuity — not for the income drop itself.

Domain D: institutional and geographic deprivation. Included only as a small, capped-weight supplementary factor — never a primary input, and never on its own sufficient to move an applicant meaningfully on the scale. Some real disadvantage — chronic under-resourcing of a specific school or region — isn't fully captured by household-level data alone, which is the justification for including it at all. But given POLAR's weakness as a proxy, its weight ceiling has to be set low enough that even a maximally scored geographic factor cannot substitute for the individual-level domains above it.

New categories of disadvantage: an evidentiary gate, not a legislative grant. Two additions were considered and are named explicitly here rather than silently excluded: a religion-linked factor, and a sexual-orientation-linked factor. Neither is resolved by fiat here, because *whether* a new category of disadvantage should be recognized at all is a legitimacy question that belongs to democratic process, not to a technical body's discretion by itself. But the process is deliberately not the one the current system uses for group recognition, because that process is exactly what Section 4 shows fails.

Recognition proceeds in three steps, each assigned to the body positioned to do it without recreating a rent.

First, evidentiary demonstration. A petitioning claim for a new domain — whether raised by a group, a legislator, or the Technical Body itself — must be supported by a measurable gap in outcomes on one or more specified metrics, and that gap must be independently reproducible by the Technical Body from its own data access, not merely asserted by the petitioner or drawn from a study the petitioner commissioned. A finding a petitioner pays for and a finding the body reproduces from population-level data are not equally trustworthy, and the design should not pretend otherwise. This closes a specific failure mode a target-based system would otherwise still be exposed to — that any sufficiently organized group can claim a metrics gap and receive floor time on the strength of the claim alone, reproducing at the recognition stage the exact copycat dynamic that followed EWS within weeks of its passage. Requiring replication, rather than mere demonstration, raises the cost of a copycat claim without closing the door to a genuine one.

Reproducibility depends on the Technical Body actually having the data to reproduce a finding, which the design does not guarantee by default. The Rohini Commission's own history is the relevant precedent: constituted in 2017 with an original twelve-week mandate, it needed disaggregated population data by sub-caste to complete its assigned analysis of OBC sub-categorization, requested budget for an all-India survey, and did not receive it — a six-year delay followed before it finally reported in 2023, still without the underlying population data it had asked for (Rohini Commission 2023). A body that depends on the same government whose policy it is evaluating for the resources to check a claim inherits exactly that vulnerability. The Technical Body's statutory funding (Section 6.5) therefore includes standing authority to commission population-level studies on any credible petition, funded from the body's own protected budget rather than a case-by-case appropriation request — otherwise the evidentiary gate is enforceable only when the government of the day chooses to make it enforceable, which defeats its purpose.

Second, legislative authorization of an outcome target, not a group entitlement. Once a gap is independently reproduced, Parliament's role is to authorize that a domain exists and to state a general outcome objective — a measurable improvement of some magnitude, on some metric, over some horizon, with the Technical Body setting the specific figures against evidence rather than Parliament by negotiation — rather than to fix a percentage share the way EWS or an OBC sub-quota does. This is the change that matters most: a percentage is an asset a beneficiary group can organize to defend once granted, which is the precise mechanism Section 4 shows entrenches every existing classification. An outcome target is not an asset in the same sense —

there is nothing to defend once the target is met, and the domain's weight is designed to converge toward zero as the gap closes, using the same convergence machinery Section 6.3 applies to weight decay within existing domains. Without this redesign, domain *creation* would have remained a legislative act, reopening at the domain boundary the exact institutional weakness this whole design exists to escape — a point developed fully in Section 6.7.

Third, administrative implementation. The Technical Body sets the specific metric, magnitude, and time horizon; owns verification and scoring within the domain exactly as it does for the four existing domains; and reports periodically on progress toward the stated target, using the same public-reporting mechanism specified in Section 6.6.

Three risks remain even with this structure, and are named rather than assumed away. First, a lowered barrier to entry is not the same as a lowered barrier to *attention* — an unorganized, dispersed population still has to get a petition onto the Technical Body's docket in the first place, and nothing here guarantees that happens for disadvantage without an organized constituency behind it. Second, the outcome target has to remain a domain-level gate on whether the domain carries weight at all, not a substitute for individual-level scoring within it; if a group-level "m percent improvement" figure starts standing in for individual verification rather than governing whether the domain stays open, group-based allocation has been reintroduced through the target's own language. Third, recognition itself — even framed as a target rather than a grant — may still trigger "if them, why not us" dynamics on the act of recognition alone, independent of whether a durable rent was created; the EWS domino effect followed within weeks of a bill passing, not specifically because of what the bill granted. None of these three risks is resolved by the evidentiary gate. They are the residual exposure that remains after it.

The first is partly answerable without abandoning the gate. The Technical Body's data-access authority (Section 6.5) shouldn't be petition-only: each term's report includes a population-level scan for gaps no one has petitioned, using the same reproducibility standard, run proactively. Petition disposition — what was received, time to hearing, reasoning for any deprioritization — is published alongside it. This doesn't equalize attention between organized and dispersed groups, but it guarantees an unpetitioned gap one occasion per term to surface, and makes suppression visible rather than silent — the same disclosure-for-representation logic already governing the no-reserved-seat design.

Applied to the two live cases: a religion-linked factor, mutually exclusive with the caste-linked Scheduled Caste factor (an asymmetry that already exists in current doctrine, since Scheduled Tribe status has no comparable religious restriction), is a coherent extension of that doctrine's own exclusivity principle — but whether an individually scored version differs enough in kind from the categorical religious quota already struck down by courts in 2010 is untested, and is exactly what the evidentiary and constitutional review has to address before implementation. A sexual-orientation-linked factor faces a different problem: no comparable administrative verification trail exists or can safely be built the way it can for every other domain, and a state-verified profile carries a disclosure and safety risk none of the other domains carry. A good-faith review would plausibly conclude this factor cannot be safely verified under the system's existing standard — a finding that would be published, not one that blocks Parliament from legislating the category anyway.

Consistency with the rest of the design. A negative replication finding — the Technical Body concluding a claimed gap does not hold up — is binding by default, on the same logic that makes the body's weighting recommendations binding by default elsewhere in this design, and is overridable only by legislative supermajority with public justification. This is deliberate: giving the Technical Body an unchecked veto over who counts as disadvantaged would create exactly the unaccountable-technocracy objection Section 5 already has to answer once for the design as a whole, and there is no reason to reopen that objection at this one juncture by treating it differently.

6.3 Setting the weights, and checking whether they are working

The specific weights within each domain are deliberately left to the administering body's statistical expertise, informed by ongoing outcome data, rather than fixed in a founding document. This carries a real cost: at launch, there is no outcomes dataset for a system that does not yet exist, so initial weights cannot be validated against outcomes the way later revisions can. They would need to be set from the best available proxy evidence — EWS's administrative record, existing caste-certificate and welfare-verification systems — and explicitly flagged as provisional. This bootstrapping problem is inherent to launching any new scoring system, and this design acknowledges it rather than assuming it away.

A further check asks whether the weights themselves are converging. A weight that genuinely tracks real disadvantage, in a system that is actually reducing that disadvantage, should trend down over time toward a floor representing "no longer predictive of extra need." A flat or rising weight signals that something needs investigating — but not, on its own, what. Two distinct failure modes produce the identical symptom: the mechanism could be broken (verification gamed, methodology miscalibrated), or the mechanism could be working correctly while the underlying disadvantage genuinely hasn't shrunk, because some forms of disadvantage — caste-based disadvantage especially — operate on a multi-generational timescale no single scoring cycle could be expected to move. The check has to distinguish the two using population-level data, not data drawn from people who already received the benefit, since their outcomes are contaminated by the effect of having received it. It also has to use domain-specific timescales: household finances are individually volatile and can move within a single business cycle for reasons that have nothing to do with whether the scoring system works, while caste-based disadvantage may show no visible movement across several review cycles by its nature.

The specific numbers this section deliberately leaves open — sub-factor weights, convergence horizons, the regional cost-of-living index for Domain B — are illustrative parameters, not load-bearing principles, and that distinction is stated explicitly rather than left implicit. The load-bearing principles are: individual rather than group eligibility; continuous rather than cliff-edge scoring; verified rather than self-declared inputs; caste retained as a factor rather than discarded; periodic recalibration against outcome data; independent administration; and transparent, appealable process. The specific figures used elsewhere in this section to illustrate the governance structure — five Appointment Panel seats, a four-year Panel term and six-year Technical Body term, a five-year cooling-off period — are each individually defended where they appear, but they are defended as reasonable instances of the principles above, not as the only numbers that could satisfy them; a different but internally consistent set of figures could serve the same design without changing what makes the architecture work.

6.4 Verification and fraud resistance, domain by domain

Caste verification routes through the existing scrutiny-committee infrastructure, which already catches some fraud, with an added second layer of audit sampling.

This layering choice needs its own honest accounting, because it is not obviously sufficient on its face. The scrutiny-committee infrastructure this design routes through is the same infrastructure whose porousness this paper documents elsewhere — fake and forged caste certificates are, as Section 6.2 already notes, a well-documented, recurring problem, and the EWS discussion at the start of this section shows a comparable existing system (self-declared, lightly cross-checked eligibility) failing in exactly the way an added audit layer is meant to catch. A second layer of sampling-based audit reduces the *rate* at which a given kind of fraud goes undetected; it does not change the *base rate* at which fraud is attempted, which is set by the underlying incentive to claim the benefit and by how much the certificate itself is trusted at first instance. Where the base infrastructure is weak enough that fraudulent certificates pass initial issuance at scale, an audit sampling layer operating downstream of that issuance step catches some fraction of what gets through — but the paper has no basis, at the design stage and before any outcome data exists, for claiming what fraction that is, and no basis for claiming it is large enough to make Domain A meaningfully more fraud-resistant than the current system rather than merely differently fraud-resistant. This is a real limit on how much of the comparative claim in Section 6.7 Domain A specifically can support, and it should be read alongside that section's own caution against overclaiming targeting accuracy: the honest position is that Domain A's verification design is a plausible improvement on the current certificate-issuance-only model, not a demonstrated one, and the state-level pilot in Section 7 (Pathway D) — which explicitly measures verification and fraud rates — is the first point at which this claim becomes checkable rather than asserted.

Financial verification relies on multi-source cross-referencing as its primary defense, since it requires falsifying convergent records across independent government databases rather than a single self-declaration. Household-circumstance verification piggybacks on whichever existing system already certifies the underlying status or benefit being used as the marker. Geographic verification carries the lowest fraud risk of the four, precisely because it is institution- or region-level rather than individually claimed — and correspondingly the lowest evidentiary value, for the same reason POLAR is weak.

Fraud resistance is not uniform across domains, and a separate risk runs alongside fraud rather than through it: verification that depends on documentation, banking history, or administrative literacy can systematically favor applicants already visible to the state, regardless of whether their claims are honest. This isn't specific to one domain, and it bears directly on how much genuine improvement individual scoring delivers over group classification — Section 6.7 returns to this. The risk isn't answered by naming it. Where PAN, GST, or land records are missing or incomplete, Domain B verification should fall back on field verification — the same audit-sampling function already assigned to Domain A — rather than defaulting an undocumented applicant to zero-evidence. The burden belongs on verification infrastructure to reach informal income, not on the applicant to produce records much of the informal workforce doesn't have. This is costlier than digital cross-referencing, which is the argument for weighting audit capacity toward informal-sector cases rather than spreading it evenly — and for reporting score and rejection rates by documentation status in the same periodic disclosures Section 6.6 already requires, so the gap is visible if it opens.

One structural property is not incidental and should be stated directly: because every domain's verified data sits in a single profile that gets re-checked at each application, a false claim in one domain is more likely to surface on re-verification than it would in a siloed, single-use system.

6.5 Who runs it: two bodies, doing different jobs

The general design principle is Majone's (2001) credible-commitment theory of delegation to non-majoritarian institutions: a government facing a temptation to intervene in its own short-term political interest can commit credibly to not doing so only by transferring authority to a body it cannot easily recapture, at some cost to direct democratic control. India's own Finance Commission already instantiates a domestic version of that principle, and the design below adopts the same separation — Parliament sets qualifications and the manner of appointment; a defined process makes the appointment; the appointees then do the substantive work — while fixing the Finance Commission's central weakness, which is that its recommendations are advisory and are sometimes simply not implemented.

The Appointment Panel is the insulation layer. It selects and can remove members of the body that does the substantive work, but does not set methodology, weights, or scoring rules itself. It has five seats, deliberately drawn from different institutional bases so no single axis can hold a majority: the executive; the political opposition, with an automatic default to the largest opposition party's leader if no recognized Leader of the Opposition exists, pre-empting a documented failure mode discussed in Section 7; the judiciary; the states collectively, through a weighted mechanism similar to the GST Council's; and an independent seat filled only from a fixed, bright-line list of eligible prior offices — not an open-ended qualifier like "eminent jurist," a term whose vagueness in comparable Indian institutional design has already been flagged as a discretion loophole. Decisions inside the panel require a supermajority of weighted votes, so no three-of-five bloc can force an appointment through alone. A separate search committee of domain specialists — not political appointees — prepares a qualification-filtered shortlist before the panel votes.

The Panel serves a **four-year term**, deliberately shorter than both the Technical Body's term and the Lok Sabha's five years — chosen so the two bodies' reconstitution dates and the electoral calendar all drift independently of one another: the Panel's four-year cycle, the Technical Body's six-year cycle, and the Lok Sabha's five-year cycle share no common factor across all three (their pairwise overlaps — four and six share a factor of two, but five shares none with either), so the three calendars do not realign on any short, predictable cycle. Because four of the five seats are tied to a currently held office rather than a personal fixed-term appointment, the Panel is **seated** at each reconstitution: whoever holds each office-linked seat at that moment serves the full four years, even if the underlying office changes hands mid-term.

The Technical Body is the substantive layer. It sets and periodically revises the deficit-point methodology and weights, owns verification and audit design, and produces the mandated public reporting discussed below. It has a chair with cross-jurisdictional public-administration experience and four members defined by qualification category: statistics and econometrics, which owns the weighting methodology and is meant to catch weak-proxy problems like POLAR's before they get built in rather than after; social stratification and caste-data specialization, requiring empirical rather than purely theoretical expertise; constitutional and legal expertise, sitting or qualified to sit on a High Court or the Supreme Court; and audit and verification expertise, owning fraud resistance and verification infrastructure. A standing secretariat separate from these five named members carries out the operational caseload.

The Technical Body serves a **six-year term**, matching the tenure India's own Comptroller and Auditor General already holds under Article 148. Up to two of the five named members may be reappointed for one further term; the remaining three or more seats must go to new appointees at each reconstitution. All members face a five-year cooling-off period afterward, long enough to outlast the government that might be positioned to reward them.

Funding and data-access independence. This is a design requirement, not a background assumption, and its absence is a documented weak point in comparable Indian institutions. TRAI's dependence on the Department of Telecommunications for its budget and staffing has repeatedly limited its ability to challenge the ministry it oversees (Telecom Regulatory Authority of India, n.d.), and the RBI's 2018 Section 7 episode — in which the government initiated, for the first time in the Reserve Bank's history, the formal consultation process under the Reserve Bank of India Act's Section 7 override provision, and Governor Urjit Patel resigned the following month amid the resulting standoff over central bank autonomy — shows even a formally insulated body remains exposed to direct pressure when a government is willing to spend the political capital (Reserve Bank of India 2018). The Technical Body's operating budget must therefore come from a statutorily protected source — a dedicated cess or fee-based mechanism, closer to SEBI's funding model than TRAI's, though SEBI's own experience is itself a qualified precedent rather than an unqualified model to follow: a 2019 amendment to the SEBI Act required SEBI to transfer three-quarters of its surplus funds to the Consolidated Fund of India and to obtain prior finance-ministry clearance for its own expenses, narrowing exactly the independence this design would want to borrow from it, and the Adani-Hindenburg episode showed separately that funding independence alone does not guarantee enforcement independence against powerful private interests (Securities and Exchange Board of India 2019). Those two qualifications do not defeat the case for fee-based funding over annual appropriation — a partially eroded independence is still meaningfully more independent than TRAI's dependence — but they mean the funding-independence requirement has to be paired with an explicit statutory bar on exactly the kind of surplus-clawback and expense-approval mechanism that narrowed SEBI's own independence after the fact, rather than assuming the SEBI model is self-sustaining. This funding explicitly includes standing authority to commission population-level studies in support of a credible petition under Section 6.2, without requiring a separate, case-by-case budget request to the government whose classification decisions are under review — the direct answer to the Rohini Commission's own experience of requesting exactly such survey funding and not receiving it.

One design choice deserves explicit defense, because it will be the first objection raised: there is no reserved caste seat on the Technical Body. Lived and field familiarity with caste-based disadvantage is a qualification criterion for the social-stratification seat, open to anyone who meets it, not a reserved allocation — a design choice defended on both political and epistemic grounds in Section 5. Legitimacy without a reserved seat comes instead from process: published methodology, a formal public consultation window before major weight revisions, and a visible grievance and appeal channel.

Removal, for members of either body, follows a two-key structure. Parliament can impeach or suspend a member by a defined vote, satisfying political accountability; actual removal is then decided by an independent judicial body assessing severity, not by a further political vote. This design avoids two failure modes visible on the record: India's own judge-removal threshold — an absolute majority plus two-thirds in both Houses — has produced zero removals even after a proven-misconduct finding against a sitting judge in 1993 (Justice V. Ramaswami), while a removal process that depended on a bare legislative majority would simply hand the body's opponents an easy capture route.

Post-retirement, a statutory bar — not a self-imposed one — prevents Technical Body members from taking any government appointment, government-linked directorship, or party position for a period long enough to outlast the political cycle they served in. Self-regulation has already been tried and abandoned in a comparable Indian context: the Supreme Court declined in 2023 to impose a cooling-off period on judges itself, calling it a legislative matter, and nothing has since been legislated.

6.6 One profile, continuously re-verified

The deficit-point system applies as a single unified score across both education admissions and public employment — one profile, not separate education- and employment-specific systems. A person's score reflects their *current*, verified circumstances, recalculated at every application point rather than drawn from a fixed lifetime allocation spent once. A one-time-use rule would reintroduce, in a different form, the exact permanence problem the system exists to solve, and would create a perverse incentive to defer claiming the larger of the two benefits, typically education, until later in life.

That design choice leaves two gaps that have to be closed separately, because they answer different questions.

Non-completion. A person who obtains a place or a position through the deficit-point system and then does not complete it — drops out, doesn't take the job, leaves early — could otherwise reappear at the next application with a full, unadjusted score, despite having consumed a scarce opportunity without converting it. The system measures *completion*, not performance: grading how well someone did once admitted would be circular, since a person admitted on a genuinely lower score because of real disadvantage may underperform relative to better-resourced peers for the same underlying reasons the score exists to account for. Completion carries no separate penalty; non-completion applies a fixed discount, set in advance, to the person's next application score — a reduction, not a reset — with a narrow, document-based exception channel for serious illness, verified institutional failure, or a formally upheld harassment or discrimination finding.

The discount is expressed as a proportional reduction, decoupling it from the point scale's eventual numeric range. Education and employment non-completion carry separate base rates, since a vacated seat and a vacated position have different displacement and training costs. A second non-completion applies an escalated multiplier, capped so it cannot compound indefinitely. And the rate itself is outcome-adjusted: at each periodic review, the administering body weighs whether non-completion among discounted applicants is falling (evidence the discount deters bad-faith claims) against whether reapplication and subsequent completion rates among previously discounted applicants are falling (evidence the discount is instead pushing out people who remain genuinely disadvantaged). If the two signals disagree, the rate stays unchanged for that term, and the report must say so. A floor keeps the discount from ever reaching zero; a ceiling keeps it from ever functioning as a disguised full reset.

Repeated successful use. A separate and independent question, easy to conflate with non-completion but analytically distinct: a person can complete every benefit they have ever claimed, in good faith, with real disadvantage each time, and still be someone for whom a scarce slot going to them meant it did not go to someone with zero prior uses. A person's first completed use in education and first completed use in employment each draw no reduction —

the ordinary single life trajectory is not repeat consumption. Reduction begins only at a second completed use within the same domain, tracked separately for education and employment. Each additional use compounds the reduction further, bounded by the same floor-and-cap logic used for the non-completion discount, and adjusted by the same two-signal outcome logic: whether repeat users are consuming a rising share of scarce slots relative to first-time applicants with comparable scores, against whether people affected by the reduction still show high verified disadvantage elsewhere in their profile — the sharper risk, since penalizing genuinely structural, multi-generational disadvantage for having "used the system too often" would reproduce the group system's permanence problem at the individual level.

The two mechanisms apply together, multiplicatively, as independent factors — which keeps them legible separately in reporting. Someone reading a periodic report should be able to see whether a reduced score reflects prior non-completion, repeated use, or both, not a single blended number.

Reporting. Each periodic report — timed to the administering body's fixed term — publishes, for each domain and each mechanism: the current rate, the signals as measured that term, and the resulting adjustment with its reasoning. This is the same transparency-in-place-of-representation principle that substitutes for a reserved seat on the Technical Body.

What this subsection and Section 6.3 have deliberately not done is set numbers. The exact sub-factors and weights within each domain, the specific discount and decay rates, the convergence horizons, the regional cost-of-living index for Domain B — all of it depends on a point scale that does not yet exist and on outcome data that cannot exist before the system launches. Fixing those numbers here, rather than leaving them to a body built specifically to revise them against evidence, would recreate exactly the failure this design is built against — a system calibrated once and defended forever, rather than one built to correct itself.

6.7 Why this particular architecture

The mechanism above is not a slogan — every documented failure mode of the systems it replaces or improves on has a specific, named answer somewhere in Section 6.1 through Section 6.6. But naming design responses to known failure modes doesn't by itself establish that this design produces a *better* allocation than group classification, only a different one. What follows states that comparative claim directly, as four conditions rather than four assumptions, because an unconditional version of any of them would be more than the mechanism above actually supports.

One clarification belongs before the four conditions, because otherwise it is easy to read more into the comparative claim than this design actually makes. None of what follows assumes that an applicant with a lower deficit-point score is less morally deserving of a scarce seat or position than one with a higher score, or that need is the only thing that could justify preferential treatment. The claim is narrower and procedural: wherever a system has already decided that preferential treatment should exist, that treatment should be targeted toward whoever has the greatest *measured* disadvantage, rather than toward whoever belongs to the classification that happened to be recognized. That is a claim about matching a given amount of preferential treatment to disadvantage more accurately, not a claim about the deeper moral basis for having preferential treatment in the first place — the two are separable, and only the first is this design's business.

The comparison is therefore evaluated on four dimensions. First, targeting accuracy: whether the rule identifies current disadvantage more closely than a categorical proxy. Second, false positives and false negatives: whether it reduces the number of substantially advantaged people receiving preferential treatment while reducing the exclusion of genuinely disadvantaged people outside the recognized category. Third, dynamic responsiveness: whether eligibility can decline as circumstances improve and rise when disadvantage emerges, rather than remaining attached to a category whose average conditions may change slowly. Fourth, political economy: whether removing the durable group boundary reduces the incentive to organize around classification itself. These are comparative hypotheses, not assumptions built into the architecture.

Individual targeting does not automatically dominate group classification on all four dimensions. It can fail through administrative neglect, underfunding, verification burden, measurement error, or systematic exclusion of people who cannot document disadvantage. Its claim to superiority is therefore conditional: individual targeting should outperform group classification when its gain in identification of current disadvantage, including reductions in false positives and false negatives, is large enough to justify its administrative cost, and when its governance is sufficiently independent to permit recalibration. It is a design that trades the concentrated capture and retrenchment problem of group-based benefits for possible administrative neglect and measurement problems; the trade is worthwhile only if the pilot evidence shows that the former losses are reduced without the latter becoming dominant.

Individual scoring targets disadvantage more accurately than group classification, where within-group variance in disadvantage is large relative to between-group variance. The Rohini Commission's own analysis of roughly 1.3 lakh central government jobs and OBC admissions found that 97 percent of those seats went to just 25 percent of OBC sub-castes, while 37 percent of OBC communities had zero representation at all — the government's own evidence that caste alone is a lossy proxy once inside a recognized category. That finding is a benefit-*capture* pattern, not itself a direct measure of underlying disadvantage, and Section 6.4 already names the risk it raises: verification that depends on documentation and administrative literacy can favor the already-visible regardless of unit of targeting. A second, independent line of evidence closes part of that gap: Zacharias and Vakulabharanam's (2009) decomposition of nationally representative household wealth survey data found that between-caste inequality accounted for only about 13 percent of overall wealth inequality in 2002–03, with substantial overlap in wealth distributions across caste categories — a meaningful share of Scheduled Caste and Scheduled Tribe households sit above the median forward-caste household, even as the group-level averages differ sharply. That is a direct measure of within-group variance relative to between-group variance, not merely a capture pattern, and it supports the pillar's premise independently of Domain A's own verification quality. This pillar holds only to the extent Domain B's multi-source verification and Domain A's audit layer actually close the capture-pattern gap rather than relocate it from sub-caste social capital to household administrative capacity — a claim about implementation quality, not something the choice of individual over group targeting guarantees on its own.

A second, distinct mechanism-level argument for individual targeting operates independently of the accuracy case above, and belongs alongside it rather than as a fifth pillar in the four-dimension comparison. Coate and Loury (1993) show that group-based preferential treatment can depress a targeted group's own investment in the skills or qualifications the treatment is meant to reward, because a group-level quota does not distinguish between members who invested and members who did not: if the marginal effort of investing is not what determines who receives the benefit, the incentive to invest weakens. This is not an accuracy argument — it holds even if group classification correctly identifies the disadvantaged population on average — and it is not a claim this paper's design otherwise makes explicitly, despite an early version of it running through Section 4's account of why classification, once made the gateway to benefit, becomes the object of contest rather than a proxy for it. A continuous, individually verified score does not eliminate the general incentive question scarcity always raises, but it does remove the specific distortion Coate and Loury identify, because the score responds to the applicant's own documented circumstances rather than to a group average that individual effort cannot move.

Accuracy is not the only thing that changes when targeting moves from groups to individuals, and a complete account has to name the other side of that trade rather than imply the design is costless to run. A large empirical literature on anti-poverty targeting across developing countries — most comprehensively synthesized by Coady, Grosh, and Hoddinott's (2004) review of 122 interventions across 47 countries — finds a consistent pattern: individually verified, continuously measured, multi-source-checked targeting is more administratively expensive to operate than a flat categorical rule, even in the cases where it is also more accurate. Verification infrastructure, periodic re-assessment, and fraud-audit capacity all cost more, in staff time and system-building, than checking a single certificate against a fixed list once. This design does not escape that tradeoff, and nothing in Section 6.1 through Section 6.6 should be read as claiming it does; the case for individual scoring is that the accuracy gain is worth the added administrative cost, not that the cost is absent.

Continuous scoring corrects two specific defects of a fixed threshold, not the fact of scarcity itself. As Section 6.1 already notes, a rank-ordered cutoff still exists wherever slots run out; continuous scoring removes the arbitrary, off-target discontinuity and disperses the gaming incentive across the scale rather than concentrating it at one visible number. That is a narrower claim than "removes the cliff-edge entirely," and the narrower version is the one this design can actually deliver.

Institutional separation resists mobilization-capture, and substitutes a different failure mode for it. A diversified Appointment Panel and an independently-mandated Technical Body raise the cost of capturing the whole system, since no single node controls both appointment and methodology. But formal insulation is not immunity: the RBI's 2018 Section 7 episode shows a government can still bring direct pressure on an insulated body when willing to spend the political capital, and TRAI's dependence on the ministry it regulates shows insulation can be hollowed out quietly instead of overridden directly (Reserve Bank of India 2018; Telecom Regulatory Authority of India, n.d.). Section 6.5's funding-independence requirement exists specifically to blunt the second risk; whether it is sufficient against the first is a design question this paper cannot resolve in advance.

A cost of the third pillar deserves the same direct treatment rather than deferral to the literature review: institutional separation blunts capture, but does not escape political vulnerability by construction — it trades one kind for another. Skocpol's (1991) finding that narrowly targeted, means-tested benefits are politically weaker than broad-based ones — because they lack a mobilized constituency with a shared interest in defending them — bears directly on the architecture proposed here. The classification-rent mechanism in Section 4 explains why a group-based system over-mobilizes to defend and expand a durable classification; the same absence of a durable, organized beneficiary class that makes an individually targeted system harder to capture also makes its budget, staffing, and institutional priority harder to defend against ordinary neglect once the political moment that created it has passed. Section 6.5's funding-independence requirement — a statutorily protected budget rather than annual appropriation — is a partial answer, aimed specifically at the case where neglect would otherwise take the form of a starved budget. It is not a full answer: a body can be adequately funded and still lose priority, attention, and appointment quality over time in a way a mobilized-constituency program does not, because no organized actor's interest is harmed by the Technical Body drifting into quiet irrelevance the way an organized actor's interest would be harmed by a group classification being narrowed. This paper does not claim to have designed around that risk; it names the risk as a real, structural cost of trading a capture-prone architecture for a neglect-prone one, not a rhetorical concession to be noted and set aside.

Self-correction holds within existing domains — and, because of Section 6.2's redesign, at the domain boundary too. The weight-convergence check in Section 6.3 tracks whether a domain's weight still reflects real disadvantage better than a fixed percentage ever could. Left as an ordinary legislative act, domain *creation* would have reopened, at exactly the point this design most needs to differ from the current system, the same institution Section 4 shows doesn't reliably revisit a settled boundary. That is why Section 6.2 replaces legislative grant with an evidentiary gate and an outcome target: it extends the same legislature-sets-goals, body-sets-means split that governs the rest of this design to the one place it had been silently missing.

The chain, and its brittleness. These four are not independent virtues bundled for convenience — each closes the failure mode the one before it would otherwise reopen. Individual targeting stays accurate only if it can be revised, which requires continuous measurement. Continuous measurement is only politically survivable if no organized bloc has a stake in what's being revised, which requires institutional separation. Separation is what makes ongoing self-correction operational rather than aspirational. That dependency makes the argument more explanatory than a list of four unrelated advantages would be — and correspondingly more exposed: if the Technical Body is captured administratively rather than politically, self-correction weakens with it, regardless of how well the first three links hold. That exposure is stated here directly, on the same principle Section 7 applies to the paper's own political-economy argument: naming a mechanism's failure mode strengthens an argument rather than weakening it.

Four conditional advantages, one dependency chain, an explicit accounting of what accuracy costs to buy, and the residual risks named at each link — including the neglect risk that institutional separation itself introduces — that is the actual shape of the comparative claim this design has to offer, not a stronger one dressed in more confident language.

6.8 What Would Falsify This Argument

The argument is intended as a testable institutional hypothesis, not only as a normative preference. Four empirical findings would materially weaken or falsify it.

Each prediction below specifies what kind of evidence would weaken or falsify it. It deliberately does not specify a numeric time window or threshold count for doing so. Fixing such a number now would repeat, at the evaluation stage, the exact mistake Section 6.3 already warns against at the design stage: treating an illustrative parameter as if it were a considered, load-bearing one. The correct evaluation window and threshold depend on the same kind of outcome data and domain-specific judgment Section 6.3 reserves for the Technical Body — mobilization dynamics, for instance, plausibly resolve faster than multi-generational persistence questions, for the same reasons Section 6.3 gives for why household-finance weights and caste-related weights should be checked on different timescales. Setting these thresholds is accordingly treated here as part of the same solidification process as the rest of the model's numeric parameters, not as a question this paper can answer in advance of the evidence base Section 7's Pathway D pilot would generate.

Prediction 1, Persistence. If established reservation categories are routinely narrowed or removed when the underlying disadvantage they were intended to address declines, the classification-rent persistence mechanism is weaker than claimed. The relevant evidence would be voluntary legislative contraction, not judicial invalidation alone, evaluated against the same evidentiary standard as Section 4's "no documented instance" finding. What counts as a sufficient interval of continued non-contraction to treat the prediction as holding, rather than merely not-yet-falsified, is left to be set once outcome or pilot data exist, consistent with Section 6.3.

Prediction 2, Mobilization. If group-based reservation produces no greater political mobilization around classification boundaries than comparable individual-targeted programs, the claim that classification creates a distinct political rent is overstated. As an operational matter, this claim is weakened if a pilot or comparable individually targeted program generates organized contestation of eligibility criteria or weights at a rate comparable to or exceeding what Section 4 documents around EWS, Maratha, Patidar, or Jat classification. The comparison window and what counts as "comparable" mobilization intensity are evaluative judgments for the body assessing pilot data, not fixed here.

Prediction 3, Targeting. If an individual deficit score performs systematically worse than group classification at identifying disadvantage, after accounting for verification, administrative cost, and measurement error, the paper's comparative claim fails. This is the one prediction among the four that does not turn primarily on a time threshold, and is correspondingly the most directly testable from a single well-designed pilot: it fails if pilot or comparative data show the deficit-point score performing at parity with or worse than the categorical baseline on the four Section 6.7 dimensions, net of verification and administrative overhead, rather than merely trading one error type for another at unchanged net accuracy.

Prediction 4, Reclassification. If individual systems rapidly generate new durable group identities and collective claims around eligibility, the proposed political-economy advantage of individual targeting is substantially reduced. As an operational matter, this claim is weakened if the system generates new, durable, organized constituencies built around specific score inputs or weight categories — sustained advocacy for a particular weight or domain, as distinct from one-off petition activity through the Section 6.2 evidentiary-gate process, which the design anticipates and is not itself evidence of reclassification. How long a constituency must persist to count as "durable" rather than transient is, again, a threshold for the administering body to set against evidence, not one this paper fixes in advance.

These tests also clarify what this paper is not claiming. A group category can remain morally or constitutionally legitimate even if the classification-rent mechanism is weak, and an individual score can be administratively feasible without being normatively preferable. The empirical question is whether the proposed change in the allocation rule produces the predicted differences in targeting, persistence, mobilization, and political durability.

---

7. Why It Won't Happen

Sections 5 and 6 argued for what the system should be and why the goal is legitimate to pursue. This section turns to a harder question: whether there is any realistic route from here to there.

An epistemic caveat that would ordinarily belong at the end of an argument like this one is stated here, at the start instead, because it should govern the tone of everything that follows rather than arrive as a late qualification: absence of precedent is not proof of impossibility. Every documented instance identified in this paper of large-scale reservation change has been expansion rather than contraction or structural redesign, and this section's finding is that no currently visible pathway changes that pattern on any near-term or medium-term timeline — a structural, evidence-backed finding about *current* conditions, not a claim that the underlying political economy is fixed for all time.

7.1 The core problem

The governance design proposed above requires constitutional-amendment-tier entrenchment for its central insulating body, so that its composition and removal process cannot be casually rewritten by an ordinary parliamentary majority. But Section 4 already showed what legislatures do with exactly this kind of opening: EWS passed both houses within days, by near-unanimous margins, precisely because opposing it cost more than supporting it. No legislative rollback of any reservation expansion appears anywhere in the evidence this paper has reviewed as a counterexample. A pathway that requires Parliament to propose and ratify a self-binding constitutional amendment is, on this paper's own evidence, asking for exactly the thing that evidence says does not happen.

Four candidate pathways are evaluated against this constraint. Three are currently unavailable or structurally insufficient for the purpose of implementing the proposal at the scale its own argument requires. One — litigation — is real, but far slower and narrower than it first appears, and does not deliver the proposal as designed even on a long timeline. None adds up to a viable near-term or medium-term route to adoption.

7.2 Pathway A: national legislation with constitutional entrenchment

The maximalist route — full replacement of the group-based system nationally, with the insulating body entrenched at constitutional-amendment level — fails for a second reason beyond the base-rate problem above. The 99th Constitutional Amendment, which created the National Judicial Appointments Commission in 2014, passed near-unanimously in both houses and was ratified by the required number of states — in other words, it cleared the exact bar this pathway would need to clear. The Supreme Court struck it down in 2015 anyway, on basic-structure grounds, holding that it compromised judicial independence.

The finding this produces is a real gap in the reasoning behind the governance design proposed in Section 6: constitutional-tier entrenchment does not remove exposure to judicial veto. It only relocates the veto point from an ordinary statute, strikeable on ordinary grounds, to a validly passed and ratified amendment that can still be voided on basic-structure grounds. Whether a technical body's methodology-and-weighting function is as exposed to that argument as judicial appointments were is not obviously true — the two are different kinds of institutional questions

— but "not obviously exposed" is not the same as "confirmed safe," and nothing in the evidence gathered here resolves that question either way.

A comparison with amendments that did succeed sharpens the point. The Twenty-Sixth Amendment, abolishing the privy purses of former princely rulers, succeeded despite an earlier Supreme Court setback; the GST Council succeeded and has held. Both had conditions this proposal does not share — a small, politically unorganized beneficiary group in the first case, mutual center-state fiscal gain in the second. Reservation beneficiaries are neither. Pathway A faces two independent barriers — legislative self-binding and post-passage judicial exposure even if self-binding somehow occurred.

7.3 Pathway B: ordinary national legislation, Lokpal-style

A second route would pass the technical body and its appointment mechanism through ordinary national legislation, avoiding Pathway A's specific failure by not requiring a constitutional amendment — the model being the Lokpal, India's anti-corruption ombudsman.

The Lokpal's own history downgrades this option rather than recommending it. It took ten failed legislative attempts across roughly forty-five years before finally passing in 2013. Even after passing, the first Lokpal chairperson was not appointed until March 2019 — over a five-year gap from the Act's passage, caused by exactly the kind of Leader-of-Opposition vacancy problem the Appointment Panel design in Section 6 defaults around from the outset. Even now, functioning, the Lokpal rejects roughly ninety percent of complaints on procedural grounds, and most state-level equivalents remain weak or unappointed.

This is not a case for abandoning the ordinary-statute route altogether — the default-to-largest-opposition-party design is a direct, evidenced response to a documented real-world failure, not speculative over-engineering. But the pathway itself remains realistically slow, appointment-stall-prone, and has a track record of being weakened in implementation even after formal passage.

7.4 Pathway C: litigation-driven reform

Courts are the only actor in the evidence reviewed here shown to have repeatedly constrained reservation's growth — the fifty percent ceiling (*Indra Sawhney v. Union of India* 1992), and the Maratha, Bihar, and Telangana interventions (*Patil v. Chief Minister* 2021; Patna High Court 2024; Telangana High Court 2025). This pathway is actually two structurally different tracks that get conflated under the general heading of "courts constrain reservation," and separating them changes the assessment substantially.

The first track (C1) is courts enforcing an existing rule against a new violation — which is what the Maratha, Bihar, and Telangana interventions, and the eventual striking-down of the NJAC amendment, actually are. This track is genuinely fast: Telangana's stay arrived within days of filing, because no new constitutional question was being decided, only an existing one applied.

It is fast precisely because it is narrow. It polices the boundary of the current system; it does not touch the system's underlying classification logic.

The second track (C2) is what the proposal actually needs: a ruling that group classification itself is an impermissible or disfavored proxy for need, opening the door to individual verification as a constitutionally required or at least permitted alternative. That is a substantial question of constitutional interpretation, which under Article 145(3) requires a bench of at least five judges — a fundamentally slower process. The Vidhi Centre's Constitution Bench Time Tracker recorded, as of February 2025, twenty-eight main matters and two hundred thirty-eight connected cases sitting with no constitution bench even constituted yet, averaging 6.4 years of pendency before a bench is formed (Vidhi Centre for Legal Policy 2025), before the case is even heard, let alone decided. Some matters have waited over a decade; the oldest, since 1986. There is no objective or binding listing mechanism; bench formation rests at the sole discretion of the sitting Chief Justice. The electoral bonds case is a useful benchmark for how this plays out even under real political salience: challenged in 2018, decided in 2024 — six years, and that case moved faster than the tracker's average. *Indra Sawhney* (1992) itself, the one case in the evidence base that did establish new doctrine via a nine-judge bench, was decided in roughly two years from the major filings of 1990–92, but under enormous and unusual political urgency that is not a realistic baseline to expect again. Even a successful C2 ruling would most plausibly establish a constraint — group proxies are constitutionally disfavored — rather than mandate this proposal's specific architecture, which returns the analysis to Pathway B's appointment-stall problem, not past it.

7.5 Pathway D: the state-level pilot as empirical bridge

A fourth option applies deficit-point scoring within an existing quota category through ordinary state legislation — for instance, how a state allocates its own share of an existing sub-quota — rather than dismantling or replacing any group's existing entitlement. Because no group's existing classification or share is being removed, there is no natural bloc with a reason to mobilize against it, the same dynamic that produced the Patidar, Jat, and Gujjar-Meena conflicts described in Section 4.

Read alongside Pathways A through C, Pathway D is best understood as an empirical test of the institutional design, not merely as an implementation workaround. It is the one place this paper's proposal can meet the real world before the harder political or judicial conditions are met. The pilot should therefore be designed as a pre-specified evaluation. It should measure predictive accuracy; false-positive and false-negative rates against independently measured indicators of disadvantage; administrative burden and processing time; verification and fraud rates; score stability across repeated applications; geographic effects; whether caste-related disadvantage is being captured rather than washed out by economic variables; whether the proposed weights predict observed disadvantage; applicant experience and appeal outcomes; and the political response relative to a comparable group-based allocation process.

The pilot can establish administrative and measurement feasibility without resolving the national constitutional or political problem of replacing existing entitlements. It can tell us whether the score can be computed, verified, audited, revised, and understood at scale; whether it performs better or worse than the categorical baseline on the four comparative dimensions in Section 6.7; and which failure modes emerge in practice. It cannot establish that Parliament would later surrender an existing group entitlement, nor can it establish that a constitutionally insulated national body would be legally permissible. Its proper contribution is narrower and more valuable: it converts a proposed architecture from a largely theoretical design into a falsifiable institutional experiment. That is a genuine contribution to the mechanics of the proposal, and it is deliberately not an answer to the escalation problem.

7.6 How the pathways relate

Track C1 does not touch the classification system either. Track C2, the only track that could, runs on a discretionary, multi-year-to-decade timeline that a state pilot's operational evidence does little to accelerate, since constitution-bench scheduling is calendar- and discretion-constrained, not evidence-constrained. The relationship between litigation and a state pilot is much thinner than sequential reinforcement: at best, a pilot supplies one input — an operational feasibility record, and outcome data a later constitutional challenge's briefing could cite — not momentum toward a faster ruling.

7.7 The finding

No pathway evaluated here presently delivers the proposal as designed on a near-term or medium-term timeline; the routes examined depend on political or judicial conditions that do not currently exist and for which the paper identifies no present trigger. Three of the four routes are currently unavailable or structurally insufficient for this specific purpose; the fourth is real only in a narrow, boundary-preserving form that does not reach the mechanism this paper identifies as problematic, while its more ambitious form encounters a discretionary bottleneck at a different institution. This is a structural, evidence-backed finding of current non-implementability, not a temporary shortage of the right strategy — and, as stated at the top of this section, not a claim that it is permanent.

---

8. What the Finding Means

The current system has a persistent expansion bias: the evidence reviewed here shows repeated expansion with little evidence of voluntary contraction, alongside an incentive structure that rewards mobilization over need. The alternative laid out in Section 6 is structurally sound in the opposite direction — bounded, self-correcting, accountable to evidence rather than to whichever group can organize hardest. And the same evidence that makes the first claim credible makes the second claim highly unlikely to be enacted under current conditions. That is not a coincidence. It is the same structure producing both outcomes: a system shaped by mobilized majorities faces strong incentives against voting to constrain those same constituencies. Section 7 is not really a separate argument from Section 4. It is Section 4 applied to its own solution.

This is a current structural incompatibility, not merely a problem of timing or advocacy. The system that would need to pass this reform is the same system the reform exists to correct. Asking it to do so is asking an institution to act against the exact incentive that defines it.

That does not make the diagnosis or the design wrong. A correct account of what is broken, and a correct account of what would fix it, do not become incorrect because the institution capable of enacting either happens to be the institution that benefits from neither existing. Section 5's own framework predicts this outcome rather than being undermined by it: contested distributive policy changes slowly, unevenly, and mostly through incremental boundary-setting rather than wholesale redesign. This structure also has real precedent in serious political philosophy — arguing for what a just system would look like while separately and honestly assessing why the existing political economy won't produce it is the basic shape of "ideal theory" work, most familiar from Rawls's (1971) distinction between ideal and non-ideal theory, though this paper's non-ideal section is far more empirically loaded than that tradition typically is. The value of ideal-theory work has never depended on the ideal being immediately achievable; it depends on the diagnosis and design being sound enough to serve as a standard against which the real system's shortfalls can be measured, and against which future, currently unforeseeable openings could be evaluated if they arise.

The point of building the alternative in Section 6 in such detail was not to submit a bill; it was to show that "individual need instead of group classification" is a real, workable system and not just a slogan — so that if the conditions that have always preceded reservation's rare moments of rapid change ever arise again, as they did for EWS, there is a design already sitting ready rather than one that has to be invented under pressure.

8.1 Contribution

This paper makes two contributions, stated here in summary form consistent with the calibrated claim already set out in the Abstract and defended against the closest prior work in Section 2. First, it states India's reservation-expansion pattern as a formal, mechanism-based diagnosis — the classification-rent mechanism named in Section 4 — rather than leaving it as the editorial observation it has generally remained in Indian political commentary. Second, building on that diagnosis, it develops a fully specified institutional alternative: a continuous, multi-domain deficit-point system administered by a two-body governance structure designed with direct reference to the documented failure modes of India's own insulated institutions — mitigating several of those modes without closing all of them, per Section 6.5 and Section 6.7 — with a falsifiable comparative claim (Section 6.7) rather than an unconditional assertion of superiority. The paper's contribution is institutional-design and political-economy work, not conceptual originality: the core move of scoring disadvantage on multiple factors rather than caste alone has circulated since at least Deshpande and Yadav (2006), and this paper's distinct claim is to have given that circulating idea governance mechanics, fraud-resistant verification, an evidentiary gate for recognizing new disadvantage categories, and an honest adoptability diagnosis that no prior version of the idea attempts.

8.2 Scope limitations

Four limitations bound what this paper's findings can support, and are stated together here rather than left scattered across the sections where each first arose. First, the case-selection logic set out in Section 3.1 treats India as a most-likely case for the classification-rent mechanism precisely because its scale, duration, and political salience make the mechanism easiest to observe; the paper's evidence for how the mechanism operates in practice is accordingly drawn from one case; the paper does not by itself establish how the same dynamic would play out in newer, smaller, or less politically entrenched quota systems. Second, Section 3.2 scopes out Malaysia's bumiputera policy and Brazil's racial quota system, both arguably closer analogues to Indian reservation on the underlying policy question than the UK comparator this paper actually draws on; a comparative treatment across those cases would be a different paper. Third, Section 6.7 names, rather than closes, a real tension the paper's own literature review surfaces: Skocpol's (1991) finding that narrowly targeted, means-tested benefits are politically weaker than broad-based ones, because they lack a mobilized constituency to defend them, predicts that the individually targeted design proposed in Section 6 may be more vulnerable to future neglect and underfunding than the group-based system it would replace, not less. Section 6.5's funding-independence requirement answers the narrow case of a starved budget; it does not answer the broader risk of a well-funded body losing priority and attention over time with no organized constituency to notice or object. This paper treats that as a real, unresolved cost of the proposed architecture, not a rhetorical concession. Fourth, the evidentiary base throughout is legal and documentary case analysis plus secondary quantitative literature, not original data collection (Section 3.3); every quantitative claim drawn from a secondary source is only as reliable as that source.

8.3 Future research

The comparative claim in Section 6.7 is stated in falsifiable form specifically so it can be tested rather than merely asserted, and three lines of future work follow directly from that structure. The most immediate is the state-level pilot described in Section 7 (Pathway D): a pre-specified evaluation of deficit-point scoring within an existing sub-quota category would generate the first real data on targeting accuracy, false-positive and false-negative rates, administrative cost, and political response relative to a comparable group-based baseline — the inputs the four predictions in Section 6.7 require and that this paper, as an institutional-design argument rather than an empirical study, does not itself supply. Second, testing Prediction 4 (reclassification) requires longitudinal tracking of whether an individually targeted system generates new durable group identities and collective claims around eligibility over time, which by definition cannot be assessed until such a system, even a pilot-scale one, has been running long enough to observe; this is a multi-year research question, not one this paper's evidentiary base can resolve in advance. Third, the comparative institutional-design gap this paper identifies relative to Malaysia and Brazil (Section 3.2) is itself a research opportunity: a systematic comparison of insulated-body governance attempts across group-based affirmative-action regimes, extending the delegation-theory analysis in Section 2 beyond the Indian institutions this paper draws on, would test whether the capture-resistant architecture proposed in Section 6.5 generalizes beyond the single domestic institutional record it is built against.

Nothing here should be read as a claim of imminent feasibility. No trigger for this change currently exists, and none is visible on the horizon. But "no trigger now" is a narrower and more accurate claim than "impossible": the system that resists reform today is not fixed for all time, only resistant under the conditions that currently hold. What this paper offers is not a near-term implementation plan. It is a standard for judging the current system and a design available for adoption if the conditions that currently block it change.

---

9. References

9.1 Legal Instruments and Cases

Association for Democratic Reforms v. Union of India. 2024. (2024) 5 SCC 1; 2024 INSC 113 [Electoral Bonds case].

Constitution (First Amendment) Act, 1951 (India).

Constitution (Twenty-Sixth Amendment) Act, 1971 (India).

Constitution (One Hundred and Third Amendment) Act, 2019 (India) [Economically Weaker Sections quota].

Constitution of India, arts. 15(4), 16(4), 145(3), 148, 243D, 279A, 334, 338B.

State of Punjab v. Davinder Singh. 2024. 2024 INSC 562 [seven-judge bench decided 1 August 2024, overruling *E.V. Chinnaiah v. State of Andhra Pradesh* (2004)].

Dr. Jaishri Laxmanrao Patil v. Chief Minister. 2021 [Maratha reservation judgment].

Government of Andhra Pradesh, G.O. Ms. No. 33 (2005), reservation scheme for Muslims, struck down by the Andhra Pradesh High Court, 2010.

Haryana creamy-layer income-only definition (2016), struck down 2021 (*Pichra Warg Kalyan Mahasabha, Haryana v. State of Haryana*, judgment delivered 24 August 2021).

Indra Sawhney v. Union of India. 1992. Supp (3) SCC 217.

Janhit Abhiyan v. Union of India. 2022 [Economically Weaker Sections case].

Kesavananda Bharati v. State of Kerala. 1973. 4 SCC 225.

Lokpal and Lokayuktas Act, 2013 (India).

NALSA v. Union of India. 2014. 5 SCC 438.

Patna High Court. 2024. Ruling against Bihar's attempted 65 percent reservation.

Supreme Court Advocates-on-Record Association v. Union of India. 2016. 5 SCC 1 (National Judicial Appointments Commission case; judgment delivered 16 October 2015).

Supreme Court of India. 2025. Order dismissing the State of Telangana's Special Leave Petition against the Telangana High Court's 9 October 2025 stay, 16 October 2025 (Vikram Nath and Sandeep Mehta, JJ.).

Telangana High Court. 2025. Interim stay of Government Order No. 9 (and consequential GOs No. 41 and 42), 9 October 2025, Backward Classes local-body reservation.

V. Ramaswami impeachment proceedings, 1993 (India).

9.2 Secondary Sources

Bhaskar, Anurag. 2022. "The Myth of the Ten-Year Limit on Reservations and Dr Ambedkar's Stance." *CASTE: A Global Journal on Social Exclusion* 3 (1).

Chandrachud, D.Y. 2021. Oral observations from the bench, Supreme Court of India, 21 October 2021, questioning the central government on the Economically Weaker Sections ₹8 lakh income cutoff and whether rural/urban differences in purchasing power had been considered. Reported in Satya Prakash, "SC Questions Centre on Rs 8 Lakh Income Norm for EWS," *The Tribune*, 21 October 2021; and "'Can't Pull Rs 8 Lakh Limit for EWS Quota Out of Thin Air': Supreme Court to Centre," *The Wire*, 23 October 2021.

Coady, David, Margaret Grosh, and John Hoddinott. 2004. *Targeting of Transfers in Developing Countries: Review of Lessons and Experience.* Washington, DC: World Bank.

Coate, Stephen, and Glenn C. Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83 (5): 1220–1240.

Counterview. 2026. "The EWS Paradox: Who Counts as Economically Weak in India Today?" counterview.net, June 2026, analyzing the EWS certification process's reliance on self-declared income and tax records, the absence of a robust verification mechanism, and the finding that under current eligibility criteria nearly 98 percent of India's population could potentially qualify for EWS benefits.

Deshpande, Ashwini, and Rajesh Ramachandran. 2024. "Caste Disparities and Affirmative Action in India." *Oxford Review of Economic Policy* 40 (3): 630–640. https://doi.org/10.1093/oxrep/grae033

Deshpande, Ashwini, and Yogendra Yadav. 2006. "Redesigning Affirmative Action: Castes and Benefits in Higher Education." *Economic and Political Weekly* 41 (24).

Galanter, Marc. 1984. *Competing Equalities: Law and the Backward Classes in India.* Berkeley: University of California Press.

Indian Masterminds. 2026. "UPSC EWS Quota Under Scanner: Investigation Finds Many Successful Candidates Had Access to Elite Education and Costly Coaching." indianmasterminds.com, 20 June 2026, drawing on an underlying investigation by *The Indian Express* into the backgrounds of all 104 candidates selected under the Economically Weaker Sections quota in the UPSC Civil Services Examination 2025.

Jaffrelot, Christophe. 2003. *India's Silent Revolution: The Rise of the Low Castes in North Indian Politics.* New York: Columbia University Press.

Jerrim, John. 2021. "Measuring Disadvantage." Sutton Trust, May 2021.

Madison, James. 1787. "Federalist No. 10." In *The Federalist Papers.*

Majone, Giandomenico. 2001. "Non-Majoritarian Institutions and the Limits of Democratic Governance: A Political Transaction-Cost Approach." *Journal of Institutional and Theoretical Economics* 157 (1): 57–78. [The credible-commitment argument is developed earlier in Majone's 1996 EUI working paper "Temporal Consistency and Policy Credibility: When Democracies Need Non-Majoritarian Institutions"; the 2001 journal article is the more citable, peer-reviewed statement of the same theory and is the version cited in text.]

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups.* Cambridge, MA: Harvard University Press.

Pierson, Paul. 1994. *Dismantling the Welfare State? Reagan, Thatcher, and the Politics of Retrenchment.* Cambridge: Cambridge University Press.

Rawls, John. 1971. *A Theory of Justice.* Cambridge, MA: Harvard University Press.

Reserve Bank of India. 2018. Government of India, consultation initiated under Section 7 of the Reserve Bank of India Act, 1934, October–November 2018, and the resignation of Governor Urjit Patel, 10 December 2018. Reported contemporaneously in Somesh Jha, "Government vs RBI: All You Need to Know about Section 7 of the RBI Act," *Business Standard*, 31 October 2018; and "RBI Governor Urjit Patel Steps Down," *Business Today*, 10 December 2018.

Rohini Commission. 2023. Report of the Commission to Examine Sub-Categorisation of Other Backward Classes. New Delhi: Government of India.

Sahu, Sangeeta. 2025. "Beyond Caste: A Multi-Factor Point-Based Reservation Model for India." *International Journal for Research in Applied Science & Engineering Technology (IJRASET)* 13 (4).

Sachar Committee. 2006. *Social, Economic and Educational Status of the Muslim Community of India: A Report.* New Delhi: Prime Minister's High Level Committee, Government of India.

Securities and Exchange Board of India. 2019. Finance Act, 2019, sec. 175, amending sec. 14 of the Securities and Exchange Board of India Act, 1992, requiring transfer of 75 percent of SEBI's annual surplus to the Consolidated Fund of India and prior finance-ministry approval of capital expenditure. Documented in Shrimi Choudhary and Somesh Jha, "Sebi Chief Questions Budget Plan for Transfer of Surplus Funds to Govt," *Business Standard*, 18 July 2019. Note: a pending Securities Markets Code, 2025 (introduced in the Lok Sabha 18 December 2025) would repeal the SEBI Act and re-enact a similar surplus-transfer provision under Clause 124. As of this note, the Code has not been enacted, but the Parliamentary Standing Committee on Finance has completed clause-by-clause examination and issued its report (June 2026), so passage may be closer than a "still with committee" framing suggests. The underlying constraint described here is expected to persist under the new Code, but this should be re-checked against the Code's final enacted text before the paper is finalized.

Skocpol, Theda. 1991. "Targeting within Universalism: Politically Viable Policies to Combat Poverty in the United States." In *The Urban Underclass*, edited by Christopher Jencks and Paul E. Peterson, 411–436. Washington, DC: Brookings Institution.

Sutton Trust. See Jerrim (2021).

Telecom Regulatory Authority of India. Budgetary and staffing dependence on the Department of Telecommunications, documented in Aditya Chunduru, "TRAI, Try Again: India's Toothless Telecom Regulator Fights for More Powers," *The Ken*, 30 August 2021 (reporting that TRAI depends on the Department of Telecommunications for budget approval and final hiring decisions, leaving three of five leadership roles vacant at the time of reporting). Note: the Department of Telecommunications notified the Telecom Regulatory Authority of India (Recruitment and Conditions of Service of Officers and Other Employees) Rules, 2025 on 18 December 2025, replacing the framework that had governed TRAI staffing since 2002 and establishing a formal Selection Committee process for direct recruitment, promotion, and deputation. This likely supersedes the specific "final hiring decisions" mechanism described in the 2021 reporting. However, TRAI's own vacancy postings continued advertising Secretary- and Advisor-level roles through at least mid-2026, so the underlying point — that TRAI runs structurally short-staffed at the leadership tier — appears to still hold even if the mechanism has changed; the specific "three of five" figure has not been independently reconfirmed for the current period and should be checked against TRAI's current leadership roster before the paper is finalized.

Teltumbde, Anand. 2024. "Sub-Categorisation Verdict: India Needs a Reservation Model Solving the Problem of Caste, Not Perpetuating It." *The Leaflet*, August 14, 2024. Republished by People's Union for Civil Liberties, pucl.org.

Vidhi Centre for Legal Policy. 2025. "Constitution Bench Time Tracker." JALDI Initiative, as of February 2025.

Weisskopf, Thomas E. 2004. *Affirmative Action in the United States and India: A Comparative Perspective.* London: Routledge.

Zacharias, Ajit, and Vamsi Vakulabharanam. 2009. "Caste and Wealth Inequality in India." Levy Economics Institute Working Paper No. 566.
